

Amendments to the drawings,

There are no amendments to the Drawings.

Remarks

Status of application

Claims 1-43 were examined and stand rejected in view of prior art. The claims have been amended to further clarify Applicant's invention. Reexamination and reconsideration are respectfully requested.

The invention

A system providing improved methods for data replication is described. In one embodiment, for example, a method of the present invention is described for replicating a transaction from a primary database to a replicate database while the replicate database remains available for use, the method comprises steps of: recording information about a transaction being performed at a primary database in a transaction log; synchronously copying the information about the transaction in the transaction log to a mirrored transaction log, so as to create at the replicate database an exact copy of the transaction log; generating a reconstructed transaction based on the information about the transaction copied to the mirrored transaction log; and applying the reconstructed transaction at the replicate database while the replicate database remains available for use.

In another embodiment, for example, a system of the present invention for replicating transactions from a source database to a standby database is described that comprises: a source database having a transaction log, the transaction log for recording log records for transactions performed at the source database; a mirrored transaction log for recording mirror copies of the log records for transactions performed at the source database, so as to create at the standby database an exact copy of the transaction log; a file mirroring module for synchronously replicating log records from the transaction log to the mirrored transaction log as transactions are performed at the source database; a log reader module for reading log records in the mirrored transaction log and reconstructing transactions for application at the standby database based upon log records in the mirrored transaction log; and a distribution module for applying the transactions reconstructed by the log reader module at the standby database.

In yet another embodiment, for example, a method of the present invention is

described for replicating a database operation from a first database to a second database while making the second database available for decision support purposes, the method comprises steps of: as a database operation is performed at the first database, generating at least one log record characterizing the operation; synchronously recording the at least one log record in a first log associated with the first database and a second log associated with the first log, so that the second log comprises an exact copy of the first log; and while the second database is available for decision support purposes, replicating the operation performed at the first database at the second database by performing the substeps of: constructing a replicate operation based, at least in part, on the at least one log record in the second log; and applying the replicate operation at the second database.

In another embodiment, for example, a method of the present invention is described for replicating transactions from a primary database to a replicate database while the replicate database remains available for use, the method comprises steps of: recording log records for transactions being performed at a primary database in a primary transaction log; creating a mirrored transaction log, the mirrored transaction log comprising an exact copy of the log records in the primary transaction log; generating reconstructed transactions based on the copies of the log records in the mirrored transaction log; and applying the reconstructed transactions at the replicate database while the replicate database remains available for use.

General: Non-statutory subject matter rejection

Claim 17 stands rejected under 35 U.S.C. 101 on the basis that the claimed invention is directed to non-statutory subject matter. The claim has been amended to restate the claimed subject matter in terms of a process or method of operation. Accordingly, the rejection on the basis of non-statutory subject matter is overcome.

Prior art rejections

A. Rejection under 35 U.S.C. 102(e): Shih

Claims 1-8, 10-23, 25-32, and 34-43 stand rejected under 35 U.S.C. 102(e) as being anticipated by Shih et al. (US 6,615,223 B1), hereinafter "Shih". The Examiner's

rejection of claims 1, 16, and 17 is representative:

As per claims 1, 16-17, Shih teaches a method and computer readable medium for replicating a transaction from a primary database to a replicate database while the replicate database remains available for use (Col. 9 lines 15-42), the method comprising:

"recording information about a transaction being performed at a primary database in a transaction log" at Col. 9 lines 15-25 and Fig. 3;

"synchronously copying the info about the transaction in the transaction log to a mirrored transaction log" at Col. 9 lines 28-60;

"generating a reconstructed transaction based on the information about the transaction copied to the mirrored transaction log" at Col. 10 lines 10-45;

"applying the reconstructed transaction at the replicate database while the replicate database remains available for use" at Col. 9 lines 28-42.

As will be shown below, Applicant's invention may be distinguished on a variety of grounds.

Shih does describe a sync replication system, and as a result shares some features in common with Applicant's system. However, Shih's approach has distinct disadvantages that Applicant's invention overcomes. (Applicant in fact had considered such a solution and discarded it due to lack of scalability and performance.) The discussion which follows identifies deficiencies of Shih and highlights features of Applicant's invention that overcome those deficiencies, as well as highlights where such features are recited as specific claim limitations present in Applicant's amended claims.

As a core architectural difference, the Shih system does not have an exact copy (mirror image) of the log at replicate sites, as log archiving and truncating is independent on the two systems. Consider the following teaching from Shih:

When a change request 12 is received at first replication site 2, server 10 issues change instruction 16 to implement the change request 12. The change instruction 16 takes into account the exact schema organization of the data object to be changed. Thus, the change instruction is schema-

specific, and in a heterogeneous environment cannot simply be sent to all remote replication sites to replicate the data change, since the schema and/or system configuration of the remote replication sites may be entirely different than the schema and system configuration of local replication site 2.

According to the invention, server 10 translates either change instruction 16 or change request 12 into a schema and system independent change record 20. Change record is in a generic format that is consistent and recognizable across all replication sites in the system. In the normal contemplated usage of the invention, change record 20 comprises change information that is focussed upon the specific data to be added, deleted, or modified by the change request 12, and does not contain information regarding the schema organization of the data at the originating replication site.

(Shih, at col. 5, lines 5-25, emphasis added.)

In order for Shih's system to support schema independence between primary and replicates, it performs the above quoted translation to write schema independent log records. (As an aside, commercial database systems generally do not write schema independent log records, thereby limiting customer appeal for Shih's approach.) In fact recall that Applicant's described system is one that operates in real-time (e.g., for supporting OLTP services), and therefore the additional processing required for writing schema independent log records would impact primary database performance in a manner that would be entirely unacceptable for such time-critical applications.

These differences between Applicant's claimed approach and that of Shih are not merely theoretical differences but ones that have important practical implications.

Consider the following teaching from Applicant's specification:

In operation, the file mirroring module 315 replicates the log records from the primary log 313 on the primary server 310 to the mirrored log 332 which is typically located at the same site as the standby server 330. The file mirroring module 315 is a disk mirroring technology that replicates the log record data very quickly and efficiently. In addition, in the currently preferred embodiment of the present invention the file mirroring module operates synchronously. When log records are written from the primary database 311 to the primary log 313, copies of these log records are written at the

same time to the mirrored primary log 332 at the standby location. This guards against any loss of data in the event of a problem with the primary database server.

(Applicant's specification, Paragraph 0060, emphasis added.)

Importantly, the design choice of Shih introduces a major disadvantage: if sync write fails, the Shih system cannot recover and must either refuse updates at the primary or completely re-materialize replicates.

Applicant's system, in contrast, does not perform Shih's resource-intensive approach of translating records. Applicant's claimed approach is one that preserves the original log file format and, thus, can work with any database system without changing the underlying database log writing function. This approach has a very important advantage. In the event that the primary database stops working, logical replication continues to be applied to the replicate database based on transactions in the mirrored primary database transaction log, enabling all of the database operations applied to the primary database to also be applied to the replicate database. This guarantees that no database operations are lost in the event of loss or damage to the primary database.

Applicant's independent claims have been amended to highlight these distinctions. For example, independent claim 1 has been amended to include the following claim limitation (shown in amended form):

synchronously copying the information about the transaction in the transaction log to a mirrored transaction log so as to create at the replicate database an exact copy of the transaction log;

(All of Applicant's other independent claims have been amended in an analogous manner.) As shown, the claim limitation requires the creation of an exact copy or mirror image of the transaction log at the replicate database. The Shih system, as described, does not teach or suggest this approach, and in fact at best Shih's translation approach appears to teach away from Applicant's claimed approach.

In view of the foregoing remarks (especially in light of clarifying amendments

made to the claims), it is respectfully submitted that the claims distinguish over Shih. Accordingly, the rejection under Section 102 is overcome.

B. Rejection under 35 U.S.C. 103(a): Shih and Riedel

Claims 9, 24 and 33 stand rejected under 35 U.S.C. 103(a) as being unpatentable over Shih as applied to claims 1-8, 10-23, 25-32, 34-43 above, and in view of Riedel et al. ("When Local Becomes Global: An Application Study of Data Consistency in a Network World"), hereinafter "Riedel." Here, the Examiner essentially repeats the rejection above (under Shih) but adds Riedel for the teaching of "said synchronously copying step includes replicating at a file block level."

The claims are believed to be allowable for at least the reasons stated above pertaining to Shih. To the point, Shih does not teach or suggest Applicant's claim limitations that require the creation of an exact copy or mirror image of the transaction log at the replicate database. The Examiner has cited nothing in Riedel that remedies that deficiency.

Further, the claims are believed to be allowable for the following additional reasons. To establish a prima facie case of obviousness under Section 103, the Examiner must establish (among other things) that (1) there is some suggestion or motivation, either in the references themselves or in the knowledge generally available to one of ordinary skill in the art, to modify the reference or to combine reference teachings, and that (2) there is a reasonable expectation of success. (See e.g., MPEP 2142). Applicant has invented an approach to read the mirrored log records in such a way as to guarantee consistency. This is done by controlling the primary database in such a way as to keep it from altering file blocks that have not been processed and by detecting when data is inconsistent. Internally (e.g., in the preferred embodiment), this requires a control file block loop to control the log archiving/log truncation at the primary database, which is of course absent from Shih's system. Merely strapping file block replication onto Shih's system does not reproduce Applicant's claimed approach. Importantly, the combined references provide no teaching or suggestion that would support the use of file block replication technique in Shih's system.

Shih in fact appears to instead be relying on Operating System file I/O and thus teaches away from the combination suggested by the Examiner. Furthermore, the Shih system would need to be substantially modify to incorporate the file block replication combination with Riedel (e.g., tending to such details as controlling the log archiving/log truncation at the primary database, as described for Applicant's system above), in order to create a system that could work in a manner (remotely) related to Applicant's system. Neither Shih nor Riedel provides any teaching, suggestion, or other motivation in that regard, and at best one would need to borrow heavily from the teachings of Applicant's specification in order have any reasonable expectation of success.

In view of the foregoing remarks, it is respectfully submitted that the claims distinguish over the combination of Shih and Riedel (especially in view of amendments to Applicant's base independent claims). Accordingly, the rejection under Section 103 is overcome.

Any dependent claims not explicitly discussed are believed to be allowable by virtue of dependency from Applicant's independent claims, as discussed in detail above.

Conclusion

In view of the foregoing remarks and the amendment to the claims, it is believed that all claims are now in condition for allowance. Hence, it is respectfully requested that the application be passed to issue at an early date.

If for any reason the Examiner feels that a telephone conference would in any way expedite prosecution of the subject application, the Examiner is invited to telephone the undersigned at 408 884 1507.

Respectfully submitted,

Date: May 2, 2006


✓
Signature
Valid

Digitally
signed by
John A. Smart
Date:
2006-05-03
15:43:21
-0700'

John A. Smart; Reg. No. 34,929
Attorney of Record

408 884 1507
815 572 8299 FAX